4.1 A 10Gb/s 5-Tap-DFE/4-Tap-FFE Transceiver in 90nm CMOS

M. Meghelli, S. Rylov, J. Bulzacchelli, W. Rhee, A. Rylyakov, H. Ainspan, B. Parker, M. Beakes, A. Chung, T. Beukema, P. Pepeljugoski, L. Shan, Y. Kwark, S. Gowda, D. Friedman

IBM, Yorktown Heights, NY

To support the high bandwidth requirements of many systems such as servers or data communication routers, low-power small-area I/O solutions are needed for serial chip-to-chip communications at line rates beyond 10Gb/s [1]. These I/Os must be capable of supporting low-cost package and board technologies that may introduce large signal degradation through bandwidth loss, reflection, and crosstalk.

In this paper, a 90nm CMOS 10Gb/s transceiver is presented. The efficient implementation of a DFE scheme in the receiver and of a FFE scheme in the transmitter allows NRZ data transmission and avoids the complexity and power consumption of a multi-level data-transmission design [2].

The transceiver design follows the basic architecture of the 0.13μm CMOS 6.4Gb/s SerDes core presented in [3]. The main design enhancements are related to the reduction of the DFE response time and the improvement of the timing-recovery precision. Also, a more power-efficient half-rate TX architecture is adopted. As shown in Fig. 4.1.1, the TX consists of a first multiplexing stage that retimes 4 single-ended quarter-rate data inputs and generates two differential half-rate even and odd data streams. These are shifted one UI with respect to each other then interleaved together to form the 1st tap of the FFE, and successively shifted by a UI then interleaved again together to form the 3 remaining taps. The 4 taps have maximum weights of {0.25, 1, 0.5, 0.25} with a resolution of {4, 6, 5, 4} bits respectively. The maximum main tap output amplitude is $1.2V_{ppd}$. Figure 4.1.2 shows the TX output eye diagram of a packaged part with -15% equalization on the 1st post-cursor compensating for ESD diode capacitance and the extra 4dB of losses of the package and evaluation board. A breakout test site of the TX is described in detail in [4].

The RX block diagram is shown in Fig. 4.1.3. A T-coil compensation network is used to mitigate the effect of the ESD diode capacitance on S_{11}. In order to ensure linear operation of the DFE, a VGA regulates the data swing at the slicer to about $0.6V_{ppd}$ (below 1-dB compression point). The VGA is designed to have 16dB of gain range and handle up to $1.2V_{ppd}$ data input swing. Besides ensuring that the analog front-end of the receiver has a wide linear range of operation and 5GHz or higher 3dB bandwidth, the most challenging part in the DFE design is to guarantee that the voltage at the slicer input (where weighted post-cursors, i.e., previously received data bits, are fed back and summed) has settled sufficiently before the data decision is made. If a classical full-rate DFE approach is used, the feedback-loop delay including the settling time needs to be less than one UI or 100ps at 10Gb/s. To ease this requirement and at the same time achieve lower power consumption, a half-rate clock DFE with speculative feedback on the first post-cursor and dynamic feedback on the remaining taps has been implemented (Fig. 4.1.3). The feedback loop delay is designed so that 2% settling accuracy is achieved within 2UI.

The clock-recovery circuit operates on the non-DFE equalized data signal and uses an Alexander-type half-rate phase detector. The early/late phase detector output is digitally filtered to generate increment/decrement signals that control a high-precision phase rotator. This phase rotator (Fig. 4.1.4) operates from two half-rate differential clock phases, I and Q. It switches the polarity of the I,Q phases (quadrant selection) and uses a 4b CML interpolator to achieve 16 phase positions within each quadrant. The phase interpolator uses a 15-cell current-steering DAC plus two additional fixed-current cells of half size to realize interpolation ratios varying from 0.5:15.5 to 15.5:0.5. Avoiding zero-value interpolation weights allows the rotator to step across each quadrant boundary by changing phase polarity only (no change in interpolation ratio). The 15 cells of the DAC are not uniform; instead, their relative sizing is optimized for best rotator linearity, with the largest cells being switched near the quadrant boundaries. Rotator linearity is also improved with the use of slew-rate-controlled buffers, which make the rotator inputs more sinusoidal. The rotator achieves a measured min-to-max step ratio better than 1:2.

A link demonstrator IC is implemented and packaged in a plastic BGA module to conduct various link experiments. The IC (Fig. 4.1.5) consists of two RX pairs and two TX pairs, each pair being either externally or internally clocked, and is configured through a parallel-port interface. The on-chip clock generation circuit consists of a full-rate LC-VCO-based PLL operating from 9GHz to 13.4GHz. The jitter generation is $<0.7ps_{rms}$ (f_c/1667–100MHz noise integration bandwidth) and the transfer bandwidth lies between 2 to 3MHz. It draws 30mA from an on-chip voltage regulator that generates a 1.2V low-noise supply from 1.8V. The power consumption of one TX/RX pair and one PLL is 300mW ($1.2V_{ppd}$ TX data ouput swing).

The link experiments presented in this paper are performed using the RX and TX pairs clocked by the on-chip PLLs at the nominal data rate of 10Gb/s. In a first experiment, a 16-inch Tyco legacy backplane channel with 24dB losses at 5GHz is successfully equalized using a stand-alone module mounted on a socketed evaluation board and used in a serial loop-back configuration. Evaluation board, plastic module, and coaxial cabling bring the total losses to 33.5dB (from the IC TX output back to the RX input). After the fixed transmitter FFE taps are configured for the channel and the DFE has adapted, the bathtub curve of the equalized serial data stream is measured. To that end, the DFE tap optimization loop is halted and the position of the phase rotator providing the data sampling clock (I-clock) is externally controlled. As shown in Fig. 4.1.6, the equalized-signal horizontal eye opening is 22% at 10^{-9} BER.

Finally, in another experiment, two modules directly soldered on a board are serially connected to each other through different channels. Figure 4.1.7 shows the horizontal eye openings at 10Gb/s and 10^{-9} BER for 10, 15, and 20-inch trace lengths with different via-stub configurations.

Acknowledgments:
The authors acknowledge funding support from the MPO; contract H98230-04-C-0920. They also wish to thank M. Sorna, S. Zier, P. Metty and K. Heilmann from IBM Fishkill for their important support.

References:
[1] "Common Electrical I/O (CEI) – Electrical and Jitter Interoperability Agreement for 6+ Gbps and 11+ Gbps I/O," *Optical Interconnect Forum*, CEI-02.0, Feb., 2005.
[2] J. L. Zerbe, et al., "Equalization and Clock Recovery for a 2.5-10Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2121-2130, Dec., 2003.
[3] M. Sorna, et al., "A 6.4Gb/s CMOS SerDes Core with Feedforward and Decision-Feedback Equalization," *ISSCC Dig. of Tech. Papers*, pp. 62-63, Feb., 2005.
[4] A. Rylyakov, et al., "A Low-Power 10 Gb/s Serial Link Transmitter in 90-nm CMOS,"*IEEE CSICS*, pp. 189-191, Nov., 2005.

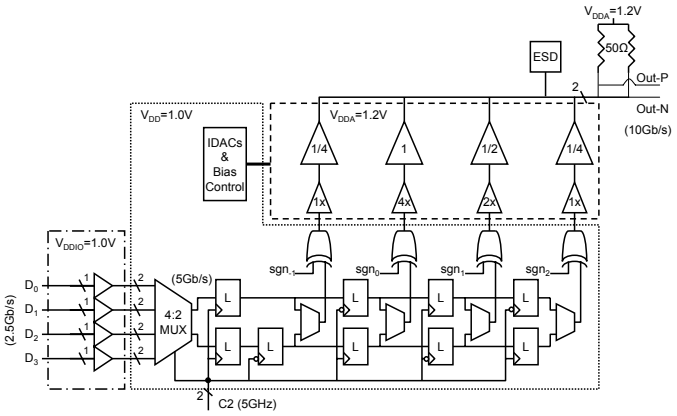

Figure 4.1.1: Transmitter block diagram.

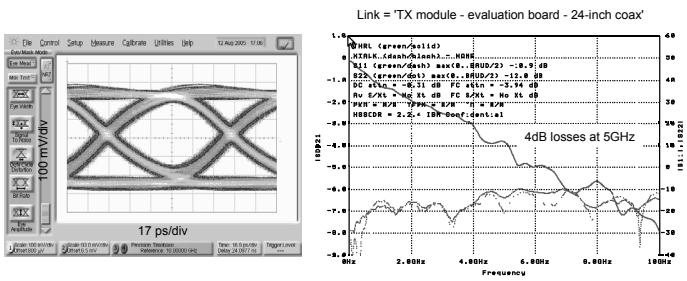

Figure 4.1.2: 10Gb/s packaged transmitter output eye diagram.

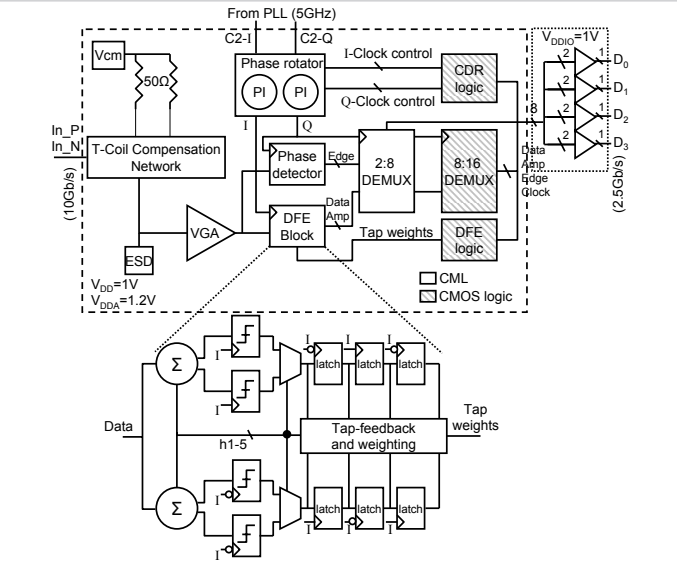

Figure 4.1.3: Receiver block diagram.

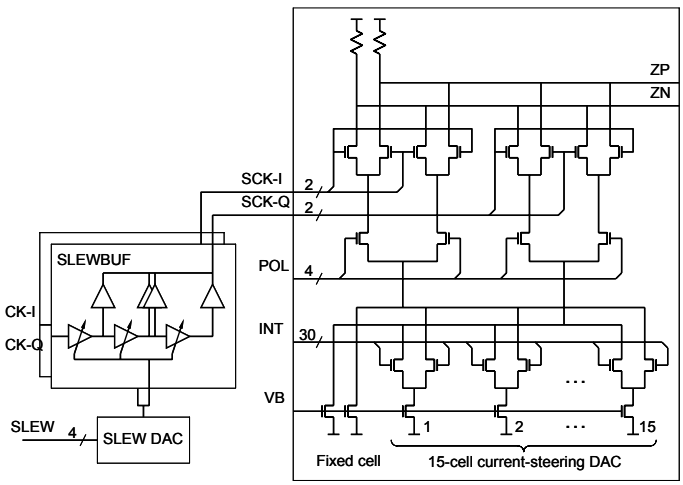

Figure 4.1.4: I,Q phase rotator schematic.

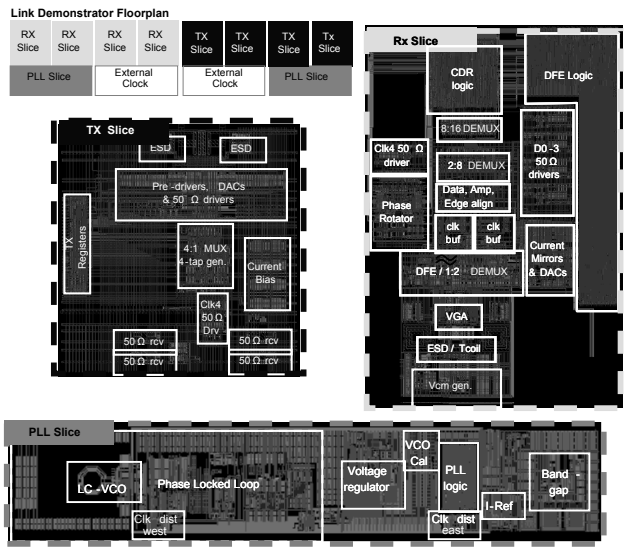

Figure 4.1.5: Link demonstrator floorplan and layout details.

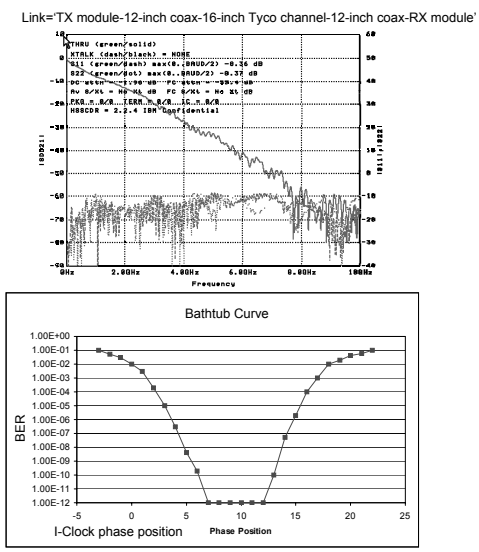

Figure 4.1.6: Equalized 16-inch tycolegacy backplane channel.

Continued on Page 639

Link: "TX module-trace-RX module" (inch)	5GHz losses	Number of vias 3.8mm via stubs / 1.8mm via stubs / 1.8mm via throughs
10 (#1)	12dB	2 / 0 / 0
10 (#2)	10dB	0 / 2 / 0
15	25dB	4 / 2 / 0
20	15dB	0 / 0 / 2

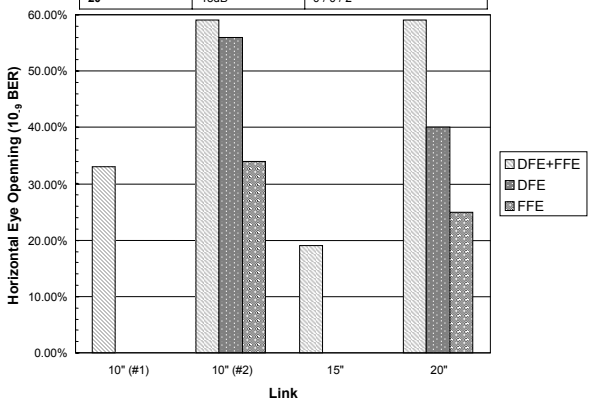

Figure 4.1.7: Chip-to-Chip link equalization experiments.

